INTERNATIONAL RESEARCH JOURNAL IN ADVANCED ENGINEERING AND TECHNOLOGY (IRJAET) E - ISSN : 2454-4752 P - ISSN : 2454-4744
VOL 2 ISSUE 1 (2016) PAGES 440-449
RECEIVED : 23/01/2016. PUBLISHED : 29/01/2016 January 29, 2016

OPTIMAL POWER MANAGEMENT SCHEME FOR CLOUD RESOURCE PROVIDERS

Dr. M. Thangamani*

Abstract- Cloud resources are provided for the users to perform computational tasks. Resources are provided by the providers with different servers. The servers are continuously running to provide resources. Energy consumption level for the servers is increased. Dynamic Voltage Frequency Scaling (DVFS) mechanism is used to adjust the voltage supply for the processors. Power supply is managed under two states. They are sleep state and active state. The power supply is reduced in the sleep state. The resource requests are consolidated with request levels. In the same way the server resource levels are also consolidated with availability details. The request and server consolidation mechanism is used in the Virtual batching technique. The Virtual batching technique is used to manage energy levels in cloud resource sharing environment. The Virtual batching mechanism is enhanced to manage memory resources. The Virtual Batching scheme is enhanced to manage resources with load balancing mechanism. The system is improved with optimization mechanism to manage relative response time. Resource levels and application requirements are integrated in the allocation process. The system is adopted to support Dynamic Random Access Memory (DRAM) and Dual in-line Memory Module (DIMM) components.

Keywords - Cloud Resources Provider, Power Management Scheme, Virtual Batching

1. INTRODUCTION

Over the last decade, the power consumption of data centers has been increasing at a rapid rate. In a report by the EPA, it is estimated that the power consumed by servers and data centers has more than doubled between the years 2000 and 2006. In 2006, it is estimated that the power consumed by servers and data centers was 61 billion kWh, which is equal to 1.5% of the total U.S. electricity consumption that year. This amounts to billion in annual electricity costs, equivalent to the power consumption costs of 5.8 million average U.S. households. Motivated by the need to reduce the power consumption of data centers, many researchers have been investigating methods to increase the energy efficiency in computing at different levels: chip, server, rack, and data center [2,5,7].

In some cases, there are physical limitations on the amount of power available for data centers. For example, Morgan Stanley is no longer able physically to get the power needed to run a new data center in Manhattan. In a survey of data centers, 31% identify power availability as a key factor limiting server deployment. The EPA report also indicates that about 50% of the power consumed in data centers is due to the infrastructure for power delivery and cooling. Therefore, minimizing the power consumed by the cooling infrastructure, can lead to significant overall power savings.

Performance states of cores provide a trade-off between the power consumed by a core and its performance. Lower P-states consume more power and provide better performance. The relationship between the performance and power consumption of the P-states is non-linear. In most cases, the lowest P-state (P0) is not the P-state with the best trade-off between performance and power consumption. P-state assignments in data centers are mainly done at the compute

 ©2016 Thangamani. M| http://www.irjaet.com

node level. In cases where the workload fluctuates, the P-state of one or more cores is increased when the node's utilization drops below a specified threshold. In a power or performance constrained data center where the workload assigned to a core is constant, the utilization of each core in a specific P-state will be close to 100% to avoid idle time. This is because we want to execute as many tasks as possible to obtain the maximum performance for a given power consumption. In this paper, we show that our technique of assigning the P-states when considering the whole data center increases the overall performance of the data center.

The power consumed by compute nodes in the data center is dissipated as heat that is removed by the CRAC units. Our approach of assigning tasks and P states to cores is thermal aware as it considers the temperature evolution effects of P-state assignments, which in turn affects the power consumed by the CRAC units. For both problems studied in this paper, we show how each assignment can be expressed as an exact optimization problem. The P-state assignment part of the problem introduces integer constraints. The integer constraints make each assignment problem not scalable with respect to the number of cores in the data center. A simple relaxation of the integer constraints may introduce additional binary constraints that make each assignment problem not scalable. To address this, we propose novel and scalable assignment techniques for both problems. Each technique involves solving a set of scalable optimization problems. Our techniques are compared against a technique that only considers putting a core in the lowest P state or turning off the core. We show that using our techniques results in notable performance improvements.

2. RELATED WORK

Server power capping enforces a power cap on a server, a number of previous approaches have been proposed in the past [6]. One possible approach is to equip each server with a feedback controller that computes the observed difference between the measured power and the power cap, and accordingly adjusts the p-state of the server using dynamic frequency and voltage scaling (DVFS). The difference is positive then DVFS is decreased, and if the difference is negative then DVFS is increased. Power budgeting determines the power cap of each server, a number of power budgeting methods have been proposed. Ghandi et al. proposed power budgeting methods for servers that execute the same workload [9]. This situation can be useful for data centers that execute transactional workloads of the same nature; they are not relevant for computing facilities that execute high-performance computing (HPC) applications. These later facilities typically have high utilizations where most of the servers are fully utilized executing a large range of workloads with potentially heterogeneous characteristics. Nathuji et al. consider the case of power budgeting for heterogeneous workloads and servers. The main proposed approach is a greedy method, where the throughput per Watt for the servers is calculated, and then servers with higher throughput per Watt are allocated more power during budgeting. Dynamic power budgeting requires the ability to estimate the impact of changing the operational power cap of the server on its performance characterizations.

A number of models have been proposed in the literature to capture the relationship between the performance and power of a single server. In early works, Rajamani et al. proposed linear models and Gandhi et al. proposed linear and cubic models [10]. The coefficients of these models are functions of the server configuration and the workload characteristics. In these works, fixed values for these coefficients were assumed irrespective of the workload characteristics. These values were obtained through prior characterization of standard benchmarks. As a result, these models are likely to show prediction errors for throughput and power in case heterogeneous applications with wide range of characteristics are executed on a cluster. In some recent works, Rountree et al. analyzed existing performance models for heterogeneous applications and proposed a linear model based on instructions per cycle and last level cache misses per cycle.

To evaluate the performance of a data center, the following metrics are usually considered: i) System normalized performance: SNP is the geometric mean of all the application normalized performance for the workloads running in the data center, where ANP is the ratio of ideal runtime to actual runtime of a workload. It is equivalent to the ratio of actual throughput to ideal throughput. Associated with ANP; ii) Slowdown norm: slowdown norm of workload i is calculated as 1=ANPi and the slowdown norm of the data center is computed by (P_i si)=N, where N is the number of workloads executing in the data center; and iii) Unfairness: Unfairness is the coefficient of variation of the ANPs for all the workloads in a data center. Thermal modeling for data center compute the minimum requirement of cooling power given a distribution of computing power, a model is needed to translate the power distribution to thermal distribution in a data center. There exist some works in thermal management in data centers. For example, Tang et al. proposed a heat cross-interference coefficient matrix based method to model the thermal distribution of data center in a fast way [8]. To quantify the efficiency of CRAC units, Moore et al. proposed a metric of coefficient of performance (CoP) [4], which enable us to build a relationship between cooling power consumption and the supply temperatures of CRAC units.

Power allocation in multi-core processors is a related problem to power budgeting in data centers [11]. Power budgeting for data centers is different in a number of ways: (1) unlike independent servers, multi-core processors do not offer power cap controllers for the individual cores; (2) workloads on a multi-core processor are likely to show memory interference issues, whereas workloads servers are relatively independent unless they explicitly communicate using message passing; (3) data centers feature air conditioning units that have to be considered during power budgeting; and (4) the interactions between computing and cooling power in data center are highly complex in nature.

3. POWER MANAGEMENT MECHANISM FOR CLOUDS

The need to provide a guaranteed level of service performance is important for data centers. This is largely due to a business model driven by strict service level agreements (SLAs) based on metrics such as response time, throughput and reserve capacity [3]. Energy demands and associated costs are increasing at an alarming rate; it is projected that data centers in the US alone will consume 100 billion kWh of energy at a cost of 7.4 billion dollars per year by 2011. This poses a dilemma for data center operators; they must satisfy new and existing service contracts while minimizing energy consumption to reduce cost and strain on power generation facilities.

Data centers generally provision based on a worst-case scenario, which leads to a low-average server utilization in modern data centers. For example, a recent estimation suggests that the utilizations of web servers are often in the 5 to 12 percent range. These underutilized servers spend a large portion of their time in an idle state. Several recent studies have shown that a server uses approximately 60 percent of its required peak power when it is idle. This over provisioning leads to large amounts of energy waste. Therefore, reducing energy waste, while guaranteeing SLA agreements, can lead to significantly reduced operating costs.

A well-known approach to addressing this problem is to transition the processor from high-power states to low power states using Dynamic Voltage and Frequency Scaling (DVFS) whenever the performance allows. This approach effectively reduces the power consumption of the computer systems when the server has a medium intensity workload [1]. The capability of this approach to reduce power consumption is limited when the server has a low-intensity workload due to two reasons. First, when the utilization of the processor is very low, the leakage power, which cannot be significantly reduced by DVFS, contributes a major portion of the power consumption. Second, many high performance processors only allow a small range of DVFS levels and even the lowest level provides a higher speed than is required for some light workloads.

To further reduce energy consumption, processors need to be put into sleep states such as Deep Sleep. In Deep Sleep, the processor is paused and consumes significantly less power. For example, the power consumption of a server with an Intel Xeon 5500 Processor may be reduced to 23 percent of its peak value when the processor is switched to the Deep Sleep state. When the processor is in Deep Sleep, the server can be configured to use Direct Memory Access (DMA) to place incoming packets into memory buffers for processing when the processor is returned to the active state, thus, avoiding harming the functionality of the hosted server applications. Therefore, to save more power for servers with light workloads, the system can perform request batching to put the processor into the Deep Sleep state when there are few incoming requests. The system delay and batch the requests when they arrive and wake the processor up when the earliest request in the batch has been kept pending for a certain batching time-out.

It is challenging to perform request batching directly on a virtualized server. Virtualization technologies such as Xen, VMware and Microsoft Hyper-V allow provisioning multiple virtual machines (VMs) onto a single physical server. However, all the VMs on a single physical server are correlated due to sharing the same physical hardware. Since different VMs may have different workloads and performance requirements, putting the processor into Deep Sleep based on the performance of one VM may affect the application performance of other VMs.

In this system proposes Virtual Batching, a novel request batching solution for virtualized enterprise servers with primarily light workloads. The solution dynamically allocates the CPU resource such that all the VMs can have approximately the same performance level relative to their allowed peak values. Based on the uniform level, the solution then determines the time length for periodically batching incoming requests and putting the processor into sleep. When the workload intensity changes from light to medium, request batching is automatically switched to DVFS to increase processor frequency for performance guarantees.

Virtual Batching is also extended to integrate with server consolidation to achieve maximized energy conservation with performance guarantees for virtualized data centers. Server consolidation can improve server utilization by consolidating VMs onto a smaller number of servers on a long time scale. Due to conservative resource profiling and various real-world constraints, servers after consolidation can still be underutilized. Virtual Batching can then be adopted to put the processors of active servers into sleep on a shorter time scale for further energy savings due to its much smaller overhead. Specifically, this system has the following contributions:

- A request batching technique is proposed in virtualized environments to achieve significant energy conservation when the server workload is light.
- Request batching and DVFS are integrated to provide energy conservation for virtualized servers when the workload varies at runtime. The system allows energy to be saved across a wide range of workload intensities.
- Two-layer control architecture relies on feedback control theory as a theoretical foundation to achieve analytical assurance of control accuracy and system stability.
- Request batching is integrated with server consolidation to achieve maximized energy conservation with performance guarantees for virtualized data centers.
- Experiments are conducted on a hardware tested with real trace files and present empirical results to demonstrate the effectiveness of the control solution to conserve energy for virtualized enterprise servers.

4. PROBLEM STATEMENT

Request batching can be conducted to group received requests into batches and put the processor into sleep between the batches. Virtual Batching is a request batching solution for virtualized servers with primarily light workloads. The system dynamically allocates CPU resources with same performance level and peak values. Server consolidation is performed to fully utilize a small number of active servers in the data center. Static and dynamic server consolidation algorithms are used to assign data centers to the request batches. Static server consolidation algorithm is used for the offline mode in data centers. Online workload variations are managed by the dynamic server consolidation algorithms. Virtual batching is integrated with pMapper (power-aware application placement framework) to assign data centers for the workloads. The following drawbacks are identified in the existing system.

- Complex virtual server sleep process
- Average relative response time is not optimized
- Energy management is tuned for the processor
- Data center load is not managed

5. OPTIMAL POWER MANAGEMENT SCHEME

The Virtual Batching scheme is enhanced to manage resources with load balancing mechanism. Fig.1 shows the scheme of optimal power management.

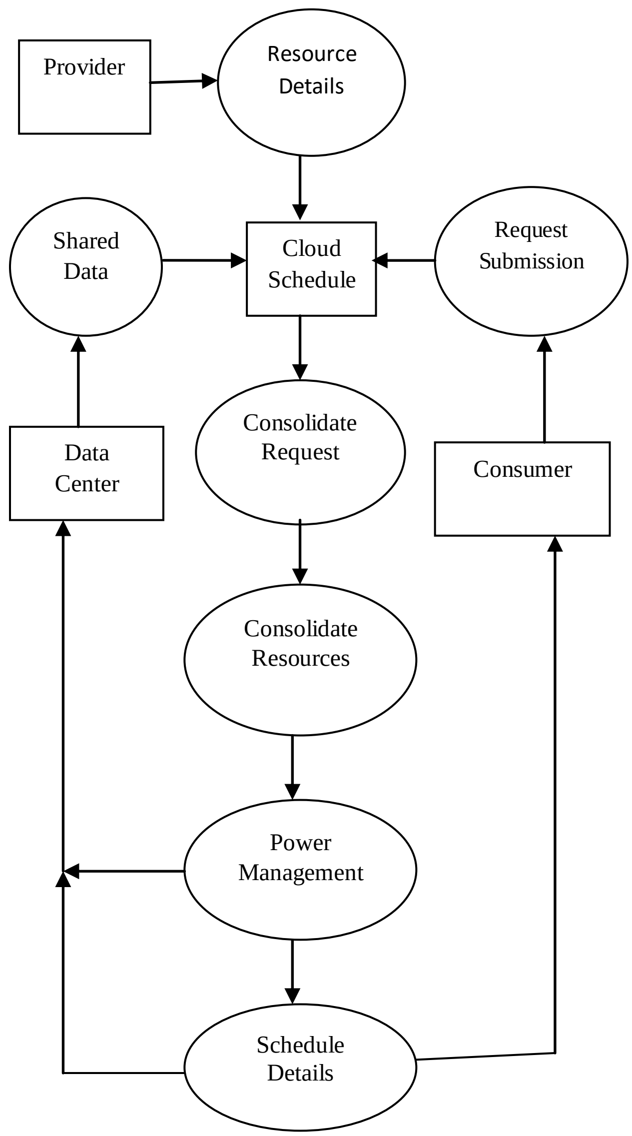

Fig.1 Optimal Power Management Scheme

The system is improved with optimization mechanism to manage relative response time. Resource levels and application requirements are integrated in the allocation process. The system is adopted to support Dynamic Random Access Memory (DRAM) and Dual in-line Memory Module (DIMM) components.

The Virtual Batching scheme is improved to manage power for computational and storage units. Request consolidation is improved with optimization techniques. Request load is distributed with different servers. The system is divided into five major modules. They are resource management, consolidation process, resource allocation process, load balancing process and power management on memory units.

Resource management module is designed to maintain the resource availability under the providers. Request and server consolidation tasks are carried out under the consolidation module. Resource allocation module handles the scheduling process. Request loads are distributed under load balancing module. Memory devices are managed with power usage levels and memory management module.

The shared resources are provided under the resource providers. Resource status is monitored for the resource providers. Total, used and available resource details are collected for each resource provider. Power supplies for the resources are controlled with DVFS support. Request and server resource levels are grouped in the consolidation process. Requests are grouped with reference to the resource requirement levels. Server consolidation is performed with resource usage and availability details. Virtual batching scheme is used for the consolidation process.

The resource allocation process is designed to schedule resources to the consumers. Request batches are used in the resource allocation process. Consolidated resources are provided for the request bundles. Scheduled resources are changed from sleep state to active state.The load balancing process is designed to distribute the request loads to the server. Server resource usage and request count details are analyzed for the load balancing process. Data center request levels are analyzed in the resource allocation process. Data center power supply is also controlled in the resource allocation process.

The DVFS scheme is used to control power supply for memory devices. The power supply for the memory devices are disconnected during the sleep state. Scheduled memory devices are tuned into active state. Dynamic Random Access Memory (DRAM) and Dual Inline Memory Module (DIMM) are managed with power supply controlling process.

6. PERFORMANCE ANALYSIS

The cloud resource scheduling scheme is designed to manage resource allocation for the tasks. The resources are provided by the resource providers. The consumer submits the requests to the scheduler. The scheduler collects thee request and resource information. Allocations of resources are carried out with the support of resource provider. The cloud energy management methods are used to reduce the power consumption in cloud resource environment. Dynamic Voltage Frequency Scaling (DVGS) mechanism is used to control the power supply in cloud resources. DVFS scheme reduces the power supply to the resource providers. Processor power supply based energy management scheme are used in the cloud resource sharing environment. The Virtual Batching Scheme (VBS) is used to manage the power consumption in cloud resource environment. Fig.2 represents the response time between the VBS and VBMS. Fig.3 and Fig.4 shows resource utilization rate analysis between VBS and VBMS and power saving analysis between VBS and VBMS.

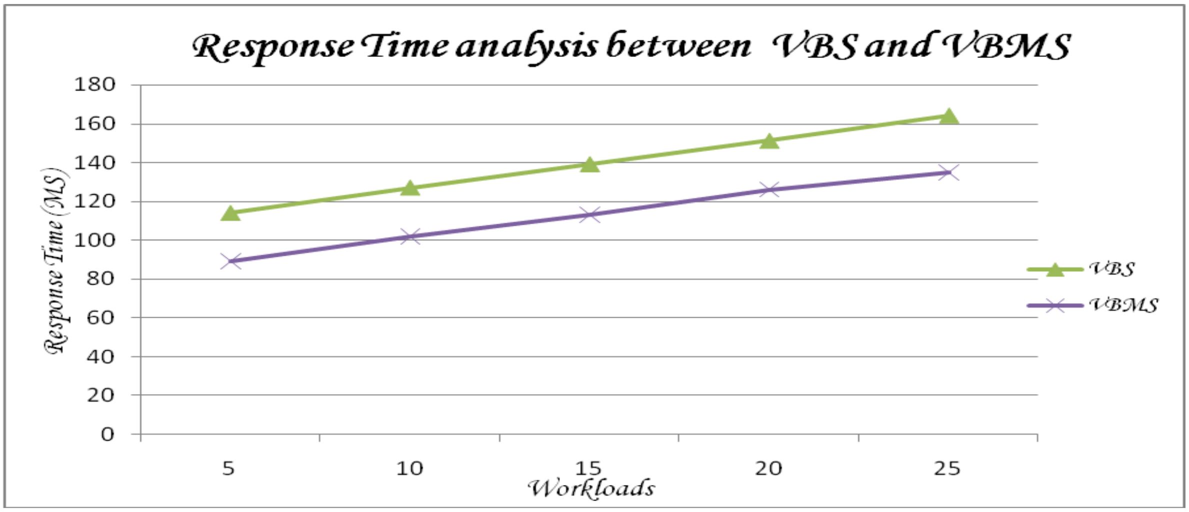

Fig.2 Response Time Analysis between VBS and VBMS

The Virtual Batching Scheme is enhanced to manage power supply in memory devices. Dynamic Random Access Memory (DRAM) and Dual Induced Memory Management (DIMM) module are managed with energy supply control mechanism. The Virtual Batching with Memory management Scheme (VBMS) is also used to control the energy usage in data centers. The energy based cloud resource scheduling system is tested with Virtual Batching Scheme (VBS) and Virtual Batching with Memory management Scheme (VBMS). The cloud resource and energy management system is tested with different resource and request levels.

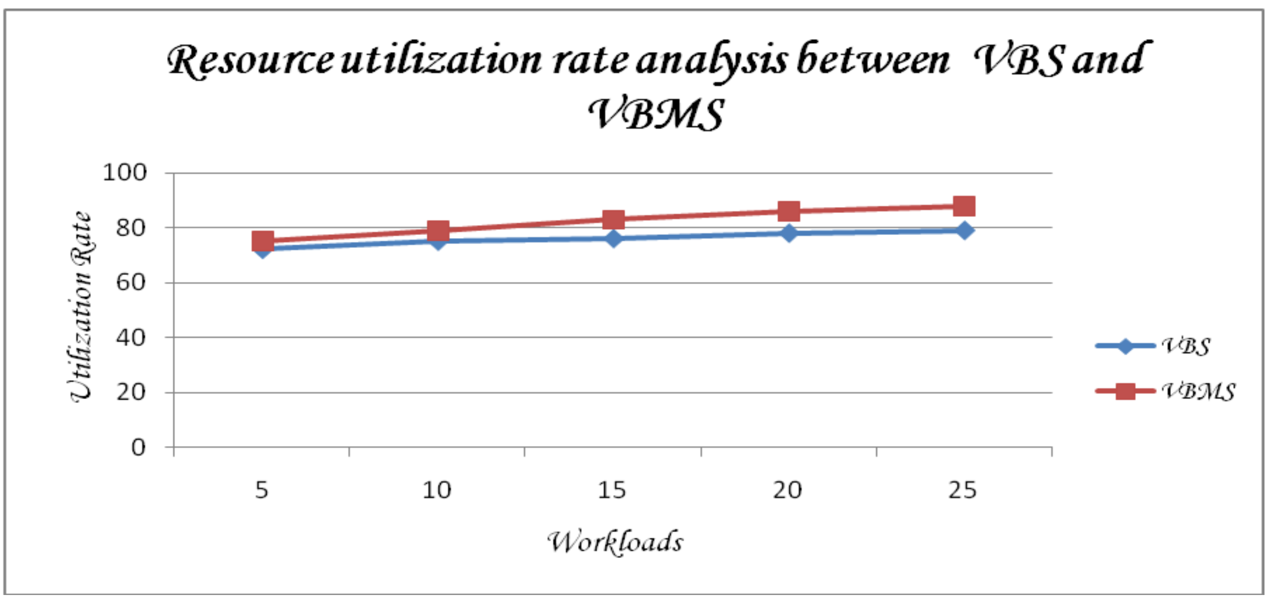

Fig. 3 Resource Utilization Rate Analysis between VBS and VBMS

©2016 Thangamani. M| http://www.irjaet.com

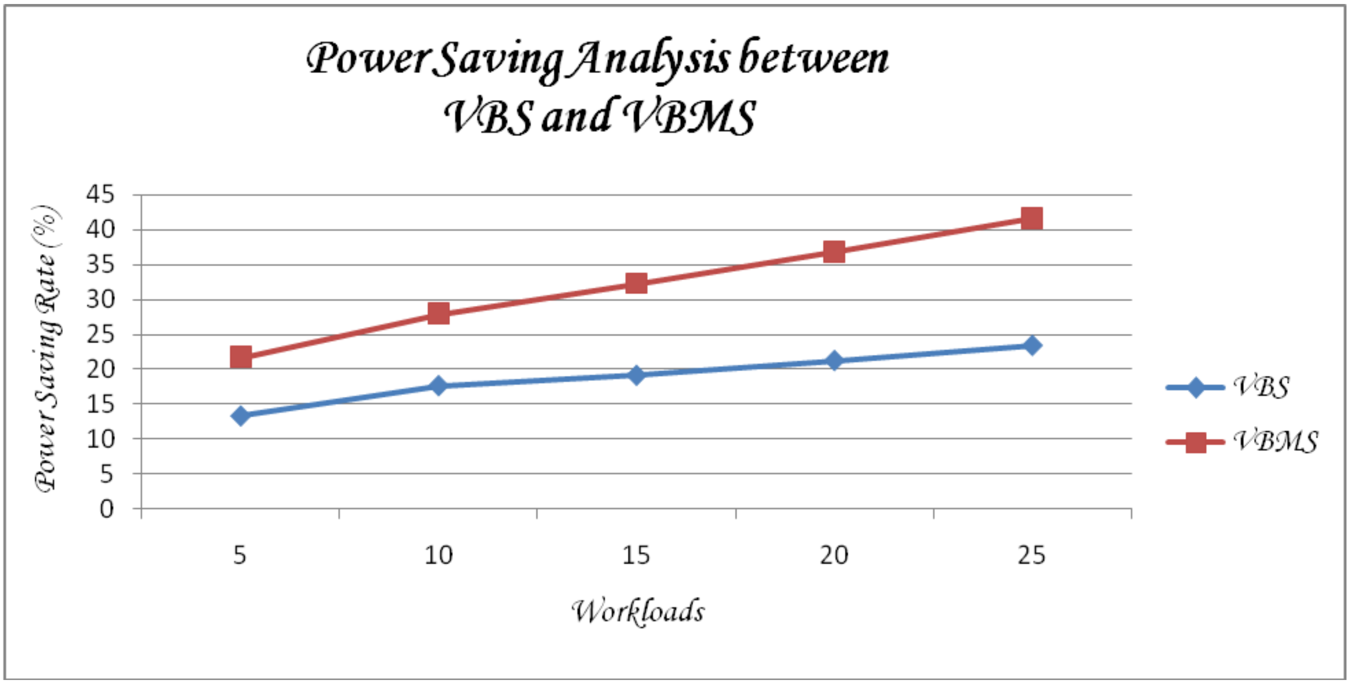

Fig.4. Power Saving Analysis between VBS and VBMS

The system uses the response time, resource utilization rate and power consumption ratio analysis for the performance process. The response time analysis for the VBS and VBMS techniques are shown in Fig.2. The VBMS reduces the average response time 20% than the VBS model. The resource utilization rate analysis for the VBS and VBMS techniques are shown in Fig. 3. The resource utilization rate in VBMS is 10% increased than the VBS mechanism. The energy saving ratio in VBMS is 35% reduced than the VBS model.

7. CONCLUSION AND FUTURE ENHANCEMENTS

The virtual batching scheme is used to reduce the energy consumption in cloud resource sharing environment. The virtual batching scheme is enhanced with memory device power control and data center power control processes. The request consolidation and server consolidation mechanisms are used to group up the requests and server resource details. The scheduling is performed with reference to the consolidated server and request details. The system is enhanced with the following features. Data security techniques can be integrated with the system. The system can be enhanced with fault tolerant mechanism to handle resource failures. The resource sharing system can be enhanced to support commercial cloud operations.

REFERENCES

[1] Abdulla M. Anthony A. Maciejewski and Howard Jay Siegel, “Power and Thermal-Aware Workload Allocation in Heterogeneous Data Centers”, IEEE Transactions on Computers, Vol. 64, No. 2, February 2015

[2] Z. Abbasi, G. Varsamopoulos, and E.K.S. Gupta, "Thermal Aware Server Provisioning and Workload Distribution for Internet Data Centers," Proc. 19th ACM Int'l Symp. High Performance Distributed Computing (HPDC'10), pp. 130-141, 2010.

[3] Yefu Wang and Xiaorui Wang, "Virtual Batching: Request Batching for Server Energy Conservation in Virtualized Data Centers", IEEE Transactions On Parallel And Distributed Systems, Vol. 24, No. 8, 2013.

[4] J. Choi, S and A. Sivasubramaniam, "Power Consumption Prediction and Power-Aware Packing in Consolidated Environments," IEEE Transaction on Computers, Vol.59, No.2, pp. 1640-1654, 2010.

[5] B.D. Young, J. Apodaca, J. Smith, B. Khemka, S. Bahirat, A. Ramirez, and Y. Zou, "Deadline and Energy Constrained Dynamic Resource Allocation in a Heterogeneous Computing Environment," J. Supercomputing, Vol. 63, No. 2, pp. 326-347, 2013.

[6] L. Parolini, N. Tolia, B. Sinopoli, and B. H. Krogh, "ACyber-Physical Systems Approach to Energy Management in Data Centers," Proc. 1st ACM/IEEE Int'l Conf. Cyber-Physical Systems, pp. 168-177, 2010.

[7] J. J. Apodaca, D. Young, L. Briceno, J. Smith, S. Pasricha, A. A. Maciejewski, H. J. Siegel, S. Bahirat, B. Khemka, A. Ramirez, and Y. Zou, "Stochastically Robust Static Resource Allocation for Energy Minimization with a Makespan Constraint in a Heterogeneous Computing Environment," The 9th IEEE International Conference on Computer Systems and Applications (AICCSA 2011), pp. 22-31, 2011.

[8] Hewlett-Packard Corporation, Intel Corporation, Microsoft Corporation, Phoenix Technologies Ltd., and Toshiba Corporation Std. (Apr. 2010). Advanced Configuration and Power Interface Specification, 2012.

[9] J.-W. Jang, M. Jeon, H.-S. Kim, H. Jo, J.-S. Kim, and S. Maeng, "Energy Reductions in Consolidated Servers through Memory-Aware Virtual Machine Scheduling," IEEE Trans. Computers, vol. 60, No. 4, 2011.

[10] A. Beloglazov, J. Abawajy, and R. Buyya, "Energy-Aware Resource Allocation Heuristics for Efficient Management of Data Centers for Cloud Computing," Future Generation Computer Systems, vol. 28, No. 5, 2012.

[11] D. Brown and C. Reams, "Toward Energy-Efficient Computing," Communications of the ACM, vol. 53, No. 3, pp. 50-58, 2010.

AUTHOR BIOGRAPHY

Dr. M. Thangamani possesses nearly 20 years of experience in research, teaching, consulting and practical application development to solve real-world business problems using analytics. Her research expertise covers Medical data mining, machine learning, cloud computing, big data, fuzzy, soft computing, ontology development, web services and open source software. She has published nearly 70 articles in refereed and indexed journals, books and book chapters and presented over 67 papers in national and international conferences in above field. She has delivered more than 60 Guest Lectures in reputed engineering colleges on various topics. She has got best paper awards from various education related social activities in India and Abroad. She has organized many self-supporting and government sponsored national conference and Workshop in the field of data mining, big data and cloud computing. She continues to actively serve the academic and research communities. She is on the editorial board and reviewing committee of leading research journals, which includes her nomination as the Associate Editor to International Journal of Entrepreneurship and Small & Medium Enterprises at Nepal and on the program committee of top international data mining and soft computing conferences in various countries. She is also seasonal reviewer in IEEE Transaction on Fuzzy System, international journal of advances in Fuzzy System and Applied mathematics and information journals. She has organizing chair and keynote speaker in international conferences in India and countries like Malaysia, Thailand and China. She has Life Membership in ISTE, Member in CSI, International Association of Engineers and Computer Scientists in China, IAENG, IRES, Athens Institute for Education and Research and Life member in Analytical Society of India. She is currently working as Assistant Professor at Kongu Engineering College at Perundurai, Erode District.